

Distributed Block Coordinate Descent for Minimizing Partially Separable Functions

Jakub Mareček, Peter Richtárik and Martin Takáč

Abstract A distributed randomized block coordinate descent method for minimizing a convex function of a huge number of variables is proposed. The complexity of the method is analyzed under the assumption that the smooth part of the objective function is partially block separable. The number of iterations required is bounded by a function of the error and the degree of separability, which extends the results in [1] to a distributed environment. Several approaches to the distribution and synchronization of the computation across a cluster of multi-core computer are described and promising computational results are provided.

1 Introduction

With the ever increasing availability of data comes the need to solve ever larger instances of problems in data science and machine learning, many of which turn out to be convex optimization problems of enormous dimensions. A single machine is often unable to store the complete data in its main memory. This suggests the need for

Jakub Mareček

IBM Research – Ireland, Dublin, Ireland, e-mail: jakub.marecek@ie.ibm.com

Peter Richtárik

School of Mathematics, University of Edinburgh, UK, e-mail: peter.richtarik@ed.ac.uk

Martin Takáč

Dept. of Industrial & Systems Engineering, Lehigh University, USA, e-mail: takac.mt@gmail.com

The first author was supported by EPSRC grant EP/I017127/1 (Mathematics for Vast Digital Resources) in 2012 and by the EU FP7 INSIGHT project (318225) subsequently. The second author was supported by EPSRC grant EP/I017127/1. The third author was supported by the Centre for Numerical Algorithms and Intelligent Software, funded by EPSRC grant EP/G036136/1 and the Scottish Funding Council.

efficient algorithms, which can benefit from distributing the data and computations across many computers.

In this paper, we study optimization problems of the form:

$$\min_{x \in \mathbf{R}^N} [F(x) := f(x) + \Omega(x)], \quad (1)$$

where f is a smooth, convex and partially block separable function, and Ω is a possibly non-smooth, convex, block separable, and “simple” extended real valued function. The technical definitions of these terms are given in Section 2.

1.1 Contributions

We propose and study the performance of a *distributed block coordinate descent method* applied to problem (1).

In our method, the blocks of coordinates are first partitioned among C computers of a cluster. Likewise, data associated with these blocks are partitioned accordingly and stored in a distributed way. In each of the subsequent iterations, each computer chooses τ blocks out of those stored locally, uniformly at random. Then, each computer computes and applies an update to the selected blocks, in parallel, out of information available to it locally. An update, which happens to be the residual in data-fitting problems, is then transmitted to other computers, which receive it either by the beginning of the next iteration or at some later time. In the former case, we denote the methods “synchronous” and we analyse them in detail. In the latter case, we denote the methods “asynchronous” and we include them for the sake of comparison in Section 7.

The main contributions of this paper are, in no particular order:

1. **Partial separability.** This is the first time such a distributed block-coordinate descent method is analyzed under the assumption that f is partially separable.
2. **New step-length.** Our method and analysis is based on an expected separable overapproximation (ESO) inequality for partially separable functions and distributed samplings in Theorem 4 in Section 4. The length of the step we take in each iteration is given by the optimum of this ESO.
3. **Iteration complexity.** We show that the iteration complexity of the method depends on the degree of block separability of f : the more separable the instance, the fewer iterations the method requires. The complexity results are stated in two theorems in Section 5 and are of the order of $O(\log(1/\varepsilon))$ for strongly convex F and $O(1/\varepsilon)$ for general convex F . At the same time, the separability also reduces the run-time per iteration.
4. **Efficient implementation.** When we replace the natural synchronous communications between computers, as analysed in Section 5, with asynchronous communication, we obtain a major speed-up in the computational performance. An efficient open-source implementation of both synchronous and asynchronous

methods is available as part of the package <http://code.google.com/p/ac-dc/>.

Our method and results are valid not only for a cluster setting, where there really are C computers which do not share any memory, and hence have to communicate by sending messages to each other, but also for computers using the Non-Uniform Memory Access (NUMA) architecture, where the memory-access time depends on the memory location relative to a processor, and accessing local memory is much faster than accessing memory elsewhere. NUMA architectures are increasingly more common in multi-processor machines.

1.2 Related work

Before we proceed, we give a brief overview of some existing literature on coordinate descent methods. For further references, we refer the reader to [2, 1, 3].

Block-coordinate descent. Block-coordinate descent is a simple iterative optimization strategy, where two subsequent iterates differ only in a single block of coordinates. In a very common special case, each block consists of a single coordinate. The choice of the block can be deterministic, e.g., cyclic ([4]), greedy ([5]), or randomized. Recent theoretical guarantees for randomized coordinate-descent algorithms can be found in [6, 7, 8, 9, 10, 11]. Coordinate descent algorithms are also closely related to coordinate relaxation, linear and non-linear Gauss-Seidel methods, subspace correction, and domain decomposition (see [12] for references). For classical references on non-randomized variants, we refer to the work of Tseng [13, 14, 15, 16].

Parallel block-coordinate descent. Clearly, one can parallelize coordinate descent by updating several blocks in parallel. The related complexity issues were studied by a number of authors. Richtárik and Takáč studied a broad class of parallel methods for the same problem we study in this paper, and introduced the concept of ESO [1]. The complexity was improved by Tappenden et al. [17]. An efficient accelerated version was introduced by Fercoq and Richtárik [3] and an inexact version was studied in [18]. An asynchronous variant was studied by Liu et al. [19]. A non-uniform sampling and a method for dealing with non-smooth functions were described in [20] and [8], respectively. Further related work can be found in [21, 22, 23, 24].

Distributed block-coordinate descent. Distributed coordinate descent was first proposed by Bertsekas and Tsitsiklis [2]. The literature on this topic was rather sparse, c.f. [25], until the research presented in this paper raised the interest, which lead to the analyses of Richtárik and Takáč [26] and Fercoq et al. [27]. These papers do not consider blocks, and specialise our results to convex functions admitting a quadratic upper bound.

In the machine-learning community, distributed algorithms have been studied for particular problems, e.g., training of support vector machines [28]. Google [29] developed a library called PSVM, where parallel row-based incomplete Cholesky fac-

torization is employed in an interior-point method. A MapReduce-based distributed algorithm for SVM was found to be effective in automatic image annotation [30]. Nevertheless, none of these papers use coordinate descent.

2 Notation and assumptions

In this section, we introduce the notation used in the rest of the paper and state our assumptions formally. We aim to keep our notation consistent with that of Nesterov [6] and Richtárik & Takáč [1].

Block structure. We decompose \mathbf{R}^N into n subspaces as follows. Let $U \in \mathbf{R}^{N \times N}$ be the $N \times N$ identity matrix and further let $U = [U_1, U_2, \dots, U_n]$ be a column decomposition of U into n submatrices, with U_i being of size $N \times N_i$, where $\sum_i N_i = N$. It is easy to observe that any vector $x \in \mathbf{R}^N$ can be written uniquely as $x = \sum_{i=1}^n U_i x^{(i)}$, where $x^{(i)} \in \mathbf{R}^{N_i}$. Moreover, $x^{(i)} = U_i^T x$. In view of the above, from now on we write $x^{(i)} := U_i^T x \in \mathbf{R}^{N_i}$, and call $x^{(i)}$ the *block* i of x .

Projection onto a set of blocks. Let us denote $\{1, 2, \dots, n\}$ by $[n]$, a set of blocks $S \subseteq [n]$, $x \in \mathbf{R}^N$, and let $x_{[S]}$ be the vector in \mathbf{R}^N whose blocks $i \in S$ are identical to those of x , but whose other blocks are zeroed out. Block-by-block, we thus have $(x_{[S]})^{(i)} = x^{(i)}$ for $i \in S$ and $(x_{[S]})^{(i)} = 0 \in \mathbf{R}^{N_i}$, otherwise. It will be more useful to us however to write

$$x_{[S]} := \sum_{i \in S} U_i x^{(i)}, \quad (2)$$

where we adopt the convention that if $S = \emptyset$, the sum is equal $0 \in \mathbf{R}^N$.

Norms. Spaces \mathbf{R}^{N_i} , $i \in [n]$, are equipped with a pair of conjugate norms: $\|t\|_{(i)}$ and $\|t\|_{(i)}^* := \max_{\|s\|_{(i)} \leq 1} \langle s, t \rangle$, $t \in \mathbf{R}^{N_i}$. For $w \in \mathbf{R}_{>0}^n$, where $\mathbf{R}_{>0}$ is a set of positive real numbers, define a pair of conjugate norms in \mathbf{R}^N by

$$\|x\|_w = \left[\sum_{i=1}^n w_i \|x^{(i)}\|_{(i)}^2 \right]^{1/2}, \quad \|y\|_w^* := \max_{\|x\|_w \leq 1} \langle y, x \rangle = \left[\sum_{i=1}^n w_i^{-1} (\|y^{(i)}\|_{(i)}^*)^2 \right]^{1/2}. \quad (3)$$

We shall assume throughout the paper that f has the following properties.

Assumption 1 (Properties of f) Function $f : \mathbf{R}^N \rightarrow \mathbf{R}$ satisfies:

1. **Partial separability.** Function f is of the form

$$f(x) = \sum_{J \in \mathcal{J}} f_J(x), \quad (4)$$

where \mathcal{J} is a collection of subsets of $[n]$ and function f_J depends on x through blocks $x^{(i)}$ for $i \in J$ only. The quantity $\omega := \max_{J \in \mathcal{J}} |J|$ is the degree of separability of f .

2. **Convexity.** Functions f_J , $J \in \mathcal{J}$ in (4) are convex.

3. **Smoothness.** The gradient of f is block Lipschitz, uniformly in x , with positive constants L_1, \dots, L_n . That is, for all $x \in \mathbf{R}^N$, $i \in [n]$ and $t \in \mathbf{R}^{N_i}$,

$$\|\nabla_i f(x + U_i t) - \nabla_i f(x)\|_{(i)}^* \leq L_i \|t\|_{(i)}, \quad (5)$$

where $\nabla_i f(x) := (\nabla f(x))^{(i)} = U_i^T \nabla f(x) \in \mathbf{R}^{N_i}$.

A few remarks are in order:

1. Note that every function f is trivially of the form (4): we can always assume that \mathcal{J} contains just the single set $J = [n]$ and let $f_J = f$. In this case we would have $\omega = n$. However, many functions appearing in applications can naturally be decomposed as a sum of a number of functions each of which depends on a small number of blocks of x only. That is, many functions have degree of separability ω that is much smaller than n .
2. Note that since f_J are convex, so is f . While it is possible to remove this assumption and provide an analysis in the non-convex case, this is beyond the scope of this paper.
3. An important consequence of (5) is the following standard inequality [31]:

$$f(x + U_i t) \leq f(x) + \langle \nabla_i f(x), t \rangle + \frac{L_i}{2} \|t\|_{(i)}^2. \quad (6)$$

Assumption 2 (Properties of Ω) We assume that $\Omega : \mathbf{R}^N \rightarrow \mathbf{R} \cup \{+\infty\}$ is (block) separable, i.e., that it can be decomposed as follows:

$$\Omega(x) = \sum_{i=1}^n \Omega_i(x^{(i)}), \quad (7)$$

where the functions $\Omega_i : \mathbf{R}^{N_i} \rightarrow \mathbf{R} \cup \{+\infty\}$ are convex and closed.

3 Distributed block coordinate descent method

In this section we describe our distributed block coordinate descent method (Algorithm 1). It is designed to solve convex optimization problems of the form (1), where the data describing the instance are so large that it is impossible to store these in memory of a single computer.

Pre-processing. Before the method is run, the set of blocks is partitioned into C sets $P^{(c)}$, $c = 1, 2, \dots, C$. Each computer “owns” one partition and will only store and update blocks of x it owns. That is, the blocks $i \in P^{(c)}$ of x are stored on and updated by computer c only. Likewise, “all data” relevant to these blocks are stored on computer c . We deal with the issues of data distribution and communication only in Section 6.

Distributed sampling of blocks. In Step 6 of Algorithm 1, each computer c chooses a random subset $Z_k^{(c)}$ of blocks from its partition $P^{(c)}$. We assume that

Algorithm Schema 1: Distributed Block Coordinate Descent

```

1 choose  $x_0 \in \mathbf{R}^N$ 
2  $k \leftarrow 0$ 
3 while termination criteria are not satisfied
4    $x_{k+1} \leftarrow x_k$ 
5   for each computer  $c \in \{1, \dots, C\}$  in parallel do
6     sample a set of coordinates  $Z_k^{(c)} \subseteq P^{(c)}$  of size  $\tau$ , uniformly at random
7     for each thread  $i \in Z_k^{(c)}$  in parallel do
8       compute an update  $h^{(i)}(x_k)$ 
9        $x_{k+1} \leftarrow x_{k+1} + U_i h^{(i)}(x_k)$ 
10   $k \leftarrow k + 1$ 

```

$|Z_k^{(c)}| = \tau$, and that it is chosen uniformly at random from all subsets of $P^{(c)}$ of cardinality τ . Moreover, we assume the choice is done independently from all history and from what the other computers do in the same iteration. Formally, we say that the set of blocks chosen by all computers in iteration k , i.e., $Z_k = \bigcup_{c=1}^C Z_k^{(c)}$, is a (C, τ) -distributed sampling.

For easier reference in the rest of the paper, we formalize the setup described above as Assumption 3 at the end of this section (where we drop the subscript k , since the samplings are independent of k).

Computing and applying block updates. In Steps 7-9, each computer c first computes and then applies updates to blocks $i \in Z_k^{(c)}$ to x_k . This is done on each computer in parallel. Hence, we have two levels of parallelism: across the nodes/computers and within each computer. The update to block i is denoted by $h^{(i)}(x_k)$ and arises as a solution of an optimization problem in the lower dimensional space \mathbf{R}^{N_i} :

$$h^{(i)}(x_k) \leftarrow \arg \min_{t \in \mathbf{R}^{N_i}} \langle \nabla_i f(x_k), t \rangle + \frac{\beta w_i}{2} \|t\|_{(i)}^2 + \Omega_i(x_k^{(i)} + t). \quad (8)$$

Our method is most effective when this optimization problem has a closed form solution, which is the case in many applications. Note that *nearly all* information that describes problem (8) for $i \in P^{(c)}$ is available at node c . In particular, $x_k^{(i)}$ is stored on c . Moreover, we can store the description of Ω_i , norm $\|\cdot\|_{(i)}$ and the pair (β, w_i) , for $i \in P^{(c)}$, on node c and only there.

Note that we did not specify yet the values of the parameters β and $w = (w_1, \dots, w_n)$. These depend on the properties of f and sampling \hat{Z} . We shall give theoretically justified formulas for these parameters in Section 4.

Communication. Finally, note that in order to find $h^{(i)}(x_k)$, each computer needs to be able to compute $\nabla_i f(x_k)$ for blocks $i \in Z_k^{(c)} \subseteq P^{(c)}$. This is the only information that an individual computer can *not* obtain from the data stored locally. We

shall describe an efficient communication protocol that allows each node to compute $\nabla_i f(x_k)$ in Section 6.

Assumption 3 (Distributed sampling) *We make the following assumptions:*

1. **Balanced partitioning.** *The set of blocks is partitioned into C groups $P^{(1)}, \dots, P^{(C)}$, each of size $s := n/C$. That is,*
 - a. $\{1, 2, \dots, n\} = \cup_{c=1}^C P^{(c)}$,
 - b. $P^{(c')} \cap P^{(c'')} = \emptyset$ for $c' \neq c''$,
 - c. $|P^{(c)}| =: s$ for all c .

2. **Sampling.** *For each $c \in \{1, \dots, C\}$, the set $\hat{Z}^{(c)}$ is a random subset of $P^{(c)}$ of size $\tau \in \{1, 2, \dots, s\}$, where each subset of size τ is chosen with equal probability.*

We refer call the random set-valued mapping $\hat{Z} := \cup_{c=1}^C \hat{Z}^{(c)}$ by the name (C, τ) -distributed sampling.

4 Expected separable overapproximation (ESO)

The following concept was first defined in [1]. It plays a key role in the complexity analysis of randomized coordinate descent methods.

Definition 1 (ESO). Let \hat{Z} be any uniform sampling, i.e., a random sampling of blocks for which $\mathbf{Prob}(i \in \hat{Z}) = \mathbf{Prob}(j \in \hat{Z})$ for all $i, j \in [n]$. We say that function f admits an ESO with respect to sampling \hat{Z} , with parameters $\beta > 0$ and $w \in \mathbf{R}_{>0}^n$, if the following inequality holds for all $x, h \in \mathbf{R}^N$:

$$\mathbf{E}[f(x + h_{[\hat{Z}]})] \leq f(x) + \frac{\mathbf{E}[|\hat{Z}|]}{n} \left(\langle \nabla f(x), h \rangle + \frac{\beta}{2} \|h\|_w^2 \right). \quad (9)$$

For simplicity, we will sometimes write $(f, \hat{Z}) \sim \text{ESO}(\beta, w)$.

In the rest of this section we derive an ESO inequality for f satisfying Assumption 1 (smooth, convex, partially separable) and for sampling \hat{Z} satisfying Assumption 3 $((C, \tau)$ -distributed sampling). This has not been done before in the literature. In particular, we give simple closed-form formulas for parameters β and w , which we shall use in Section 5 to shed light on the performance of the method.

We first need to establish an auxiliary result. We use $[n]$ to denote $\{1, 2, \dots, n\}$.

Lemma 1. *Let $\hat{Z} = \cup_{c=1}^C \hat{Z}^{(c)}$ be a (C, τ) -distributed sampling. Pick $J \subseteq [n]$ and assume that $|P^{(c)} \cap J| = \xi$ for some $\xi \geq 1$ and all c . Let $\kappa = \kappa(|\hat{Z} \cap J|, i)$ be any function that depends on $|\hat{Z} \cap J|$ and $i \in [n]$ only. Then*

$$\mathbf{E} \left[\sum_{i \in \hat{Z} \cap J} \kappa(|\hat{Z} \cap J|, i) \right] = \mathbf{E} \left[\frac{|\hat{Z} \cap J|}{C\xi} \sum_{i \in J} \kappa(|\hat{Z} \cap J|, i) \right]. \quad (10)$$

Proof. Let us denote by $J^{(c)} = J \cap P^{(c)}$, $\zeta = |\hat{Z} \cap J|$ and $\zeta^{(c)} = |\hat{Z} \cap J^{(c)}|$. Then

$$\begin{aligned}
\mathbf{E} \left[\sum_{i \in \hat{Z} \cap J} \kappa(\zeta, i) \right] &= \mathbf{E} \left[\mathbf{E} \left[\sum_{i \in \hat{Z} \cap J} \kappa(\zeta, i) \mid \zeta \right] \right] \\
&= \mathbf{E} \left[\mathbf{E} \left[\mathbf{E} \left[\sum_{i \in \hat{Z} \cap J} \kappa \left(\sum_{c=1}^C \zeta^{(c)}, i \right) \mid \zeta^{(1)}, \dots, \zeta^{(C)}, \sum_{c=1}^C \zeta^{(c)} = \zeta \right] \mid \zeta \right] \right] \\
&= \mathbf{E} \left[\mathbf{E} \left[\mathbf{E} \left[\sum_{c=1}^C \sum_{i \in \hat{Z}^{(c)} \cap J^{(c)}} \kappa(\zeta, i) \mid \zeta^{(1)}, \dots, \zeta^{(C)} \right] \mid \sum_{c=1}^C \zeta^{(c)} = \zeta \right] \right] \\
&= \mathbf{E} \left[\mathbf{E} \left[\sum_{c=1}^C \frac{\zeta^{(c)}}{\xi} \sum_{i \in J^{(c)}} \kappa(\zeta, i) \mid \sum_{c=1}^C \zeta^{(c)} = \zeta \right] \right] \\
&= \mathbf{E} \left[\sum_{c=1}^C \frac{\zeta}{\xi C} \sum_{i \in J^{(c)}} \kappa(\zeta, i) \right] = \mathbf{E} \left[\frac{\zeta}{\xi C} \sum_{i \in J} \kappa(\zeta, i) \right]. \quad \square
\end{aligned}$$

The main technical result of this paper follows. This is a generalization of a result from [1] for partially separable f and τ -nice sampling to the distributed ($c > 1$) case. Notice that for $C = 1$ we have $\xi = \omega$.

Theorem 4 (ESO). *Let f satisfy Assumption 1 and \hat{Z} satisfy Assumption 3. Let ¹ $\xi := \max\{|P^{(c)} \cap J| : c \in \{1, \dots, C\}, J \in \mathcal{J}\}$. Then (f, \hat{Z}) admits ESO with parameters β and w given by*

$$\beta = 1 + \frac{(\xi - 1)(\tau - 1)}{\max\{1, s - 1\}} + (C - 1) \frac{\xi \tau}{s}, \quad (11)$$

and $w_i = L_i$, $i = 1, 2, \dots, n$.

Proof. For fixed $x \in \mathbf{R}^N$, define $\phi(h) := f(x + h) - f(x) - \langle \nabla f(x), h \rangle$. Likewise, for all $J \in \mathcal{J}$ we define $\phi_J(h) := f_J(x + h) - f_J(x) - \langle \nabla f_J(x), h \rangle$. Note that

$$\phi(h) = \sum_{J \in \mathcal{J}} \phi_J(h). \quad (12)$$

Also note that the functions ϕ_J and ϕ are convex and minimized at $h = 0$, where they attain the value of 0. For any uniform sampling, and hence for \hat{Z} in particular, and any $a \in \mathbf{R}^N$, one has $\mathbf{E}[\langle a, h_{[\hat{Z}]}\rangle] = \frac{\mathbf{E}[\|\hat{Z}\|]}{n} \langle a, h \rangle$, and therefore

$$\mathbf{E}[\phi(h_{[\hat{Z}]})] = \mathbf{E}[f(x + h_{[\hat{Z}]})] - f(x) - \frac{\mathbf{E}[\|\hat{Z}\|]}{n} \langle \nabla f(x), h \rangle. \quad (13)$$

Because of this, and in view of (9) and the fact that as $\mathbf{E}[\|\hat{Z}\|] = C\tau$,² we only need to show that

¹ Note that $\xi \in \{\lceil \frac{\omega}{C} \rceil, \dots, \omega\}$.

² In fact, $\|\hat{Z}\| = C\tau$ with probability 1.

$$\mathbf{E}[\phi(h_{[\hat{Z}]})] \leq \frac{C\tau}{n} \frac{\beta}{2} \|h\|_w^2. \quad (14)$$

Our starting point in establishing (14) will be the observation that from (6) used with $t = h^{(i)}$ we get

$$\phi(U_i h^{(i)}) \leq \frac{L_i}{2} \|h^{(i)}\|_{(i)}^2, \quad i \in [n]. \quad (15)$$

To simplify the proof, we shall without loss of generality assume that $|P^{(c)} \cap J| = \xi$ for all $c \in \{1, 2, \dots, C\}$ and $J \in \mathcal{J}$ for some constant $\xi > 1$. This can be achieved by extending the sets $J \in \mathcal{J}$ by introducing dummy dependencies (note that the assumptions of the theorem are still satisfied after this change). For brevity, let us write $\theta_{J,\hat{Z}} := |J \cap \hat{Z}|$ and $h_{[i]} := U_i h^{(i)}$. Fixing $J \in \mathcal{J}$ and $h \in \mathbf{R}^N$, we can estimate:

$$\begin{aligned} \mathbf{E}[\phi_J(h_{[\hat{Z}]})] &\stackrel{(2)}{=} \mathbf{E} \left[\phi_J \left(\sum_{i \in \hat{Z}} h_{[i]} \right) \right] = \mathbf{E} \left[\phi_J \left(\sum_{i \in \hat{Z} \cap J} h_{[i]} \right) \right] \\ &= \mathbf{E} \left[\phi_J \left(\frac{1}{\theta_{J,\hat{Z}}} \sum_{i \in \hat{Z} \cap J} \theta_{J,\hat{Z}} h_{[i]} \right) \right] \leq \mathbf{E} \left[\frac{1}{\theta_{J,\hat{Z}}} \sum_{i \in \hat{Z} \cap J} \phi_J \left(\theta_{J,\hat{Z}} h_{[i]} \right) \right] \\ &\stackrel{(10)}{=} \mathbf{E} \left[\frac{1}{\theta_{J,\hat{Z}}} \left(\frac{\theta_{J,\hat{Z}}}{C\xi} \sum_{i \in J} \phi_J \left(\theta_{J,\hat{Z}} h_{[i]} \right) \right) \right] = \frac{1}{C\xi} \mathbf{E} \left[\sum_{i \in J} \phi_J \left(\theta_{J,\hat{Z}} h_{[i]} \right) \right] \\ &= \frac{1}{C\xi} \mathbf{E} \left[\sum_{i \in [n]} \phi_J \left(\theta_{J,\hat{Z}} h_{[i]} \right) \right]. \end{aligned} \quad (16)$$

In the second equation above we have used the assumption that ϕ_J depends on blocks $i \in J$ only. The only inequality above follows from convexity of ϕ_J . Note that this step can only be performed if the sum is over a nonempty index set, which happens precisely when $\theta_{J,\hat{Z}} \geq 1$. This technicality can be handled at the expense of introducing a heavier notation (which we shall not do here), and (16) still holds. Finally, in one of the last steps we have used (10) with $\kappa(\hat{Z} \cap J, i) \leftarrow \phi_J(\theta_{J,\hat{Z}} h_{[i]})$.

By summing up inequalities (16) for $J \in \mathcal{J}$, we get

$$\begin{aligned} \mathbf{E}[\phi(h_{[\hat{Z}]})] &\stackrel{(12)}{=} \sum_{J \in \mathcal{J}} \mathbf{E}[\phi_J(h_{[\hat{Z}]})] \stackrel{(16)}{\leq} \frac{1}{C\xi} \sum_{J \in \mathcal{J}} \mathbf{E} \left[\sum_{i \in [n]} \phi_J \left(\theta_{J,\hat{Z}} h_{[i]} \right) \right] \\ &\stackrel{(12)}{=} \frac{1}{C\xi} \mathbf{E} \left[\sum_{i \in [n]} \phi \left(\theta_{J,\hat{Z}} h_{[i]} \right) \right] \stackrel{(15)}{\leq} \frac{1}{C\xi} \mathbf{E} \left[\sum_{i \in [n]} \frac{L_i}{2} \|\theta_{J,\hat{Z}} h^{(i)}\|_{(i)}^2 \right] \\ &= \frac{1}{2C\xi} \mathbf{E} \left[\theta_{J,\hat{Z}}^2 \sum_{i \in [n]} L_i \|h^{(i)}\|_{(i)}^2 \right] \stackrel{(3)}{=} \frac{1}{2C\xi} \|h\|_w^2 \mathbf{E}[\theta_{J,\hat{Z}}^2]. \end{aligned} \quad (17)$$

We now need to compute $\mathbf{E}[\theta_{J,\hat{Z}}^2]$. Note that the random variable $\theta_{J,\hat{Z}}$ is the sum of C independent random variables $\theta_{J,\hat{Z}} = \sum_{c=1}^C \theta_{J,\hat{Z}^{(c)}}$, where $\theta_{J,\hat{Z}^{(c)}}$ has the simple law

$$\mathbf{Prob}(\theta_{J,\hat{Z}^{(c)}} = k) = \binom{\xi}{k} \binom{s-\xi}{\tau-k} / \binom{s}{\tau}.$$

We therefore get

$$\begin{aligned} \mathbf{E}[\theta_{J,\hat{Z}}^2] &= \mathbf{E}\left[\left(\sum_{c=1}^C \theta_{J,\hat{Z}^{(c)}}\right)^2\right] = C\mathbf{E}[(\theta_{J,\hat{Z}^{(c)}})^2] + C(C-1)(\mathbf{E}[\theta_{J,\hat{Z}^{(c)}}])^2 \\ &= C\frac{\xi\tau}{s} \left(1 + \frac{(\xi-1)(\tau-1)}{\max\{1, s-1\}}\right) + C(C-1) \left(\frac{\xi}{s}\tau\right)^2. \end{aligned} \quad (18)$$

It only remains to combine (17) and (18) to get (14). \square

Note that ESO inequalities have recently been used in the analysis of distributed coordinate descent methods by Richtárik and Takáč [26] and Fercoq et al. [27] However, their assumptions on f and derivation of ESO are very different and hence our results apply to a different class of functions.

5 Iteration complexity

In this section, we state two iteration complexity results for Algorithm 1. Theorem 5 deals with a non-strongly convex objective and shows that the algorithm achieves sub-linear rate of convergence $\mathcal{O}(\frac{1}{\varepsilon})$. Theorem 6 shows Algorithm 1 achieves linear convergence rate $\mathcal{O}(\log \frac{1}{\varepsilon})$ for a strongly convex objective.

However, we wish to stress that in high dimensional settings, and especially in applications where low- or medium-accuracy solutions are acceptable, the dependence of the method on ε is somewhat less important than its dependence on data size through quantities such as the dimension N and the number of blocks n , and on quantities such as the number of computers C and number of parallel updates per computer τ , which is related to the number of cores.

Notice that once the ESO is established by Theorem 4, the complexity results, Theorems 5 and 6, follow from the generic complexity results in [17] and [1], respectively.

5.1 Convex functions

Theorem 5 (Based on [17]). *Let f satisfy Assumption 1 and sampling \hat{Z} satisfy Assumption 3. Let x_k be the iterates of Algorithm 1 applied to problem (1), where parameters β and w are chosen as in Theorem 4 and the random sets Z_k are iid, following the law of \hat{Z} . Then for all $k \geq 1$,*

$$\mathbf{E}[F(x_k) - F^*] \leq \frac{n}{n + C\tau k} \left(\frac{\beta}{2} \|x_0 - x^*\|_w^2 + F(x_0) - F^* \right). \quad (19)$$

Note that the leading term in the bound decreases as the number of blocks updated in a single (parallel) iteration, $C\tau$, increases. However, notice that the parameter β also depends on C and τ . We shall investigate this phenomenon in Section 5.3 and show that the level of speed-up one gets by increasing C and/or τ (where by speed-up we mean the decrease of the upper bound established by the theorem) depends on the degree of separability ω of f . The smaller ω is, the more speed-up one obtains.

5.2 Strongly-convex functions

If we assume that F is strongly convex with respect to the norm $\|\cdot\|_w$ then the following theorem shows that $F(x_k)$ converges to F^* linearly, with high probability.

Definition 2 (Strong convexity). Function $\phi : \mathbf{R}^N \rightarrow \mathbf{R} \cup \{+\infty\}$ is strongly convex with respect to the norm $\|\cdot\|_w$ with convexity parameter $\mu_\phi(w) \geq 0$ if

$$\phi(y) \geq \phi(x) + \langle \phi'(x), y - x \rangle + \frac{\mu_\phi(w)}{2} \|y - x\|_w^2, \quad \forall x, y \in \text{dom } \phi, \quad (20)$$

where $\phi'(x)$ is any subgradient of ϕ at x .

Notice that by setting $\mu_\phi(w) = 0$, one obtains the usual notion of convexity. Strong convexity of F may come from f or Ω (or both); we write $\mu_f(w)$ (resp. $\mu_\Omega(w)$) for the (strong) convexity parameter of f (resp. Ω). It follows from (20) that if f and Ω are strongly convex, then F is strongly convex with, e.g., $\mu_F(w) \geq \mu_f(w) + \mu_\Omega(w)$.

Theorem 6 (Based on [1]). *Let us adopt the same assumptions as in Theorem 5. Moreover, assume that F is strongly convex with $\mu_f(w) + \mu_\Omega(w) > 0$. Choose initial point $x_0 \in \mathbf{R}^N$, target confidence level $0 < \rho < 1$, target accuracy level $0 < \varepsilon < F(x_0) - F^*$ and*

$$K \geq \frac{n}{C\tau} \frac{\beta + \mu_\Omega(w)}{\mu_f(w) + \mu_\Omega(w)} \log \left(\frac{F(x_0) - F^*}{\varepsilon \rho} \right). \quad (21)$$

If $\{x_k\}$ are the random points generated by Algorithm 1, then $\mathbf{Prob}(F(x_K) - F^ \leq \varepsilon) \geq 1 - \rho$.*

Notice that now both ε and ρ appear inside a logarithm. Hence, it is easy to obtain accurate solutions with high probability.

5.3 Parallelization speed-up is governed by sparsity

If we assume that $\|x_0 - x^*\|_w^2 \gg F(x_0) - F^*$, then in view of Theorem 5, the number of iterations required by our method to get an ε solution in expectation is $O(\frac{\beta}{C\tau\varepsilon})$.




Fig. 1 Speed-up gained from updating more blocks per iteration is almost linear initially, and depending on sparsity level η , may become significantly sublinear.

Hence, the smaller $\frac{\beta}{C\tau\epsilon}$ is, the fewer are the iterations required. If β were a constant independent of C and τ , one would achieve linear speed-up by increasing workload (i.e., by increasing $C\tau$). However, this is the case for $C = 1$ and $\omega = 1$ only (see Theorem 4). Let us look at the general case. If we write $\eta := \frac{\xi}{s}$ (this a measure of sparsity of the partitioned data), then

$$\begin{aligned} \frac{\beta}{C\tau} &\stackrel{(11)}{=} \frac{1 + \frac{(\xi-1)(\tau-1)}{\max\{1, s-1\}} + (C-1)\frac{\xi\tau}{s}}{C\tau} \leq \frac{1 + \frac{\xi(\tau-1)}{s} + (C-1)\frac{\xi\tau}{s}}{C\tau} \\ &= \frac{1 + \eta(\tau-1) + (C-1)\eta\tau}{C\tau} = \frac{1 + \eta(C\tau-1)}{C\tau} = \frac{1}{C\tau} + \eta \left(1 - \frac{1}{C\tau}\right). \end{aligned}$$

As expected, the first term represents linear speed-up. The second term represents a penalty for the lack of sparsity (correlations) in the data. As $C\tau$ increases, the second term becomes increasingly dominant, and hence slows the speed-up from almost linear to none. Notice that for fixed η , the ratio $\frac{\beta}{C\tau}$ as a function of $C\tau$ is decreasing and hence we always get *some* speed-up by increasing $C\tau$.

Figure 1 (left) shows the speed-up factor ($\frac{C\tau}{\beta}$; high values are good) as a function of $C\tau$ for different sparsity levels η . One can observe that sparse problems achieve almost linear speed-up even for bigger value of $C\tau$, whereas for, e.g., $\eta = 0.2$, almost linear speed-up is possible only up to $C\tau = 10$. For sparser data with $\eta = 0.01$, linear speed-up can be achieved up to $C\tau = 100$. For $\eta = 0.001$, we can use $C\tau = 10^3$. The right part of Figure 1 shows how sparsity affects speed-up for a fixed number of updates $C\tau$. Again, the break-point of almost linear speed-up is visibly present.

Similar observations in the non-distributed setting were reported in [1]. The phenomenon is not merely a by-product of our theoretical analysis; it also appears in practice.

5.4 The cost of distribution

Notice that in a certain intuitive sense, variants of Algorithm 1 are comparable, as long as each iteration updates the same number $C\tau$ of blocks. This allows us to vary C and τ , while keeping the product constant. In particular, let us consider two scenarios:

1. Consider C computers, each updating τ blocks in parallel, and
2. Consider 1 computer updating $C\tau$ blocks in each iteration in parallel.

For the sake of comparison, we assume that the underlying problem is small enough so that it can be stored on and solved by a single computer. Further, we assume that F is strongly convex, $\mu(\Omega) = 0$ and $s = \frac{n}{C} \geq 2$. Similar comparisons can be made in other settings as well, but given the page restrictions, we restrict ourselves to this case only.

In the iteration-complexity bound (21), we notice that the only difference is in the value of β . Let β_1 be the β parameter in the first situation with C computers, and β_2 be the β parameter in the second situation with 1 computer. The ratio of the complexity bounds (21) is hence equal to the ratio

$$\frac{\beta_1}{\beta_2} = \frac{(1 + \frac{(\xi-1)(\tau-1)}{s-1}) + (C-1)\frac{\xi\tau}{s}}{1 + \frac{(\omega-1)(C\tau-1)}{Cs-1}}.$$

Notice that $\frac{\omega}{C} \leq \xi \leq \omega$. The ratio β_1/β_2 is increasing in ξ . We thus obtain the following bounds:

$$\text{LB} := \frac{1 + \frac{(\omega-C)(\tau-1)}{n-C} + (C-1)\frac{\omega\tau}{n}}{1 + \frac{(\omega-1)(C\tau-1)}{n-1}} \leq \frac{\beta_1}{\beta_2} \leq \frac{1 + \frac{(\omega-1)(C\tau-C)}{n-C} + (C-1)\frac{\omega C\tau}{n}}{1 + \frac{(\omega-1)(C\tau-1)}{n-1}} =: \text{UB}.$$

Table 1 presents the values of LB and UB for various parameter choices and problem sizes. We observe that the value of β_2 is around 1. The value of β_1 depends on a particular partition, but we are sure that $\beta_1 \in [\beta_2 \cdot \text{LU}, \beta_2 \cdot \text{UB}]$. In Table 1, UB is less than 2, which means that by distributing the computation, the method will at most double the number of iterations. However, larger values of UB, albeit $\text{UB} \lesssim C$, are possible for different settings of the parameters. For a different class of functions f , an upper bound of 2 was proven in [26] and improved in [27] to the factor $1 + 1/(\tau - 1)$ whenever $\tau > 1$.

Of course, if the problem size exceeds the memory available at a single computer, the option of not distributing the data and computation may not be available. It is reassuring, though, to know that the price we pay for distributing the data and computation, in terms of the number of iterations, is bounded. Having said that, a major complication associated with any distributed method is the communication, which we discuss in the two following sections.

Table 1 Lower and upper bounds on β_1/β_2 for a selection parameters n, ω, C and τ .

| n | ω | C | τ | β_2 | LB | UB |
|--------|----------|-----|--------|-----------|-----------|-----------|
| 10^6 | 10^2 | 10 | 50 | 1.049 | 1.0000086 | 1.4279673 |
| 10^7 | 10^2 | 10 | 50 | 1.005 | 1.0000009 | 1.0446901 |
| 10^8 | 10^2 | 100 | 100 | 1.009 | 1.0000010 | 1.9801990 |

6 Two implementations

Although our algorithm and results apply to a rather broad class of functions, we focus on two important problems in statistics and machine learning in describing our computational experience, so as to highlight the finer details of the implementations.

6.1 An implementation for sparse least squares

In many statistical analyses, e.g., linear regression, one hopes to find a solution x with only a few non-zero elements, which improves interpretability. It has been recognized, however, that the inclusion of the number of non-zero elements, $\|x\|_0$, in the objective function raises the complexity of many efficiently solvable problems to NP-Hard [32, 33]. Recently, a number of randomized coordinate descent methods try to handle the ℓ_0 -norm directly [34], but only local convergence can be guaranteed. Fortunately, the inclusion of the sum of absolute values, $\|x\|_1$, provides a provably good proxy, which is also known as ℓ_1 regularization. There is a large and growing body of work on both practical solvers for non-smooth convex problems, obtained by such a regularization, and their convergence properties, when one restricts oneself to a single computer storing the complete input. Such solvers are, however, most useful in high-dimensional applications, where the size of the data sets often exceeds the capacity of random-access memory of any single computer available today.

Hence, the first implementation we present is a distributed coordinate-descent algorithm for ℓ_1 -regularized (“sparse”) least squares. The key components needed by Algorithm 1 are the computation of L_i , $\nabla_i f(x_k)$, and solving of a block-wise minimization problem. Note that $\nabla_i f(x) = \sum_{j=1}^m -A_{j,i}(y^{(j)} - A_{j,:}x)$, where $A_{j,:}$ denotes j -th row of matrix A , and $L_i = \|A_{:,i}\|_2^2$. The only difficulty is that given the data partition $\{P^{(c)}\}_{c=1}^C$, no single computer c is able to compute $\nabla_i f(x)$ for any $i \in P^{(c)}$. The reasoning follows from a simple observation: if we wanted to compute $\nabla_i f(x_k)$ for a given x_k from scratch, we would have to access all coordinates of x_k , vector y , and all non-zero elements of the input matrix A . This could be avoided by introducing an auxiliary vector $g_k := g(x_k)$ defined as

$$g_k := Ax_k - y. \quad (22)$$

Once the value of $g_k = g(x_k)$ is available, a new iterate is

$$x_{k+1} = x_k + \sum_{c=1}^C \sum_{i \in Z_k^{(c)}} U_i h^{(i)}(x_k). \quad (23)$$

and $g_{k+1} = g(x_{k+1})$ can be easily expressed as

$$g_{k+1} = g_k + \underbrace{\sum_{c=1}^C \sum_{i \in Z_k^{(c)}} A_{:i} h^{(i)}(x_k)}_{\delta g^{(c)}}. \quad (24)$$

Note that the value $\delta g^{(c)}$ can be computed on computer c as all required data are available on computer c . Subsequently, g_{k+1} can be obtained by summation and the formula for $\nabla_i f(x)$ will take the form $\nabla_i f(x) = A_{:i}^T g = \sum_{j=1}^m A_{j,i} g^{(j)}$. Once we know how to compute $\nabla_i f(x)$ and L_i , all that remains to be done is to solve the problem

$$\min_{t \in \mathbf{R}} a + bt + \frac{c}{2} t^2 + \lambda |d + t|, \quad (25)$$

where $a, b, d \in \mathbf{R}$ and $c, \lambda \in \mathbf{R}_{>0}$, which is given by a *soft-thresholding* formula $t^* = \text{sgn}(\zeta)(|\zeta| - \frac{\lambda}{c})_+ - d$, where $\zeta = d - \frac{b}{c}$.

6.2 An implementation for training support vector machines

Let us present another example implementation. The key problem in supervised machine learning is the training of classifiers. Given a matrix $A \in \mathbf{R}^{m \times N}$, a compatible vector $y \in \mathbf{R}^m$, and constant $\gamma > 0$, the goal is to find a vector $x \in \mathbf{R}^N$ which solves the following optimization problem:

$$\min_{x \in \mathbf{R}^N} F(x) := \underbrace{\gamma \|x\|_1}_{\Omega(x)} + \underbrace{\sum_{j=1}^m \mathcal{L}(x, A_{j:}, y^{(j)})}_{f(x)}, \quad (26)$$

where $A_{j:}$ again denotes j -th row of matrix A and \mathcal{L} is a loss function, such as

$$\mathcal{L}_{SL}(x, A_{j:}, y^{(j)}) := \frac{1}{2} (y^{(j)} - A_{j:}x)^2, \quad \text{square loss,} \quad (\text{SL})$$

$$\mathcal{L}_{LL}(x, A_{j:}, y^{(j)}) := \log(1 + e^{-y^{(j)} A_{j:}x}), \quad \text{logistic loss,} \quad (\text{LL})$$

$$\mathcal{L}_{HL}(x, A_{j:}, y^{(j)}) := \frac{1}{2} \max\{0, 1 - y^{(j)} A_{j:}x\}^2, \quad \text{hinge square loss.} \quad (\text{HL})$$

The input (A, y) is often referred to as the training data. Rows of matrix A represent observations of N features each and y are the corresponding classifications to train the classifier on.

Square hinge loss is a popular choice of \mathcal{L} , but is not smooth. It is well known that the dual has the form [25, 35, 36]:

$$\min_{x \in \mathbf{R}^m} F(x) := \underbrace{\frac{1}{2\lambda m^2} x^T Q x - \frac{1}{m} x^T \mathbf{1}}_{f(x)} + \underbrace{\sum_{i=1}^m \Phi_{[0,1]}(x^{(i)})}_{\Omega(x)}, \quad (\text{SVM-DUAL})$$

where $\Phi_{[0,1]}$ is the characteristic (or “indicator”) function of the interval $[0, 1]$ and $Q \in \mathbf{R}^{m \times m}$ is the Gram matrix of the data, i.e., $Q_{i,j} = y^{(i)} y^{(j)} A_i A_j^T$. If x^* is an optimal solution of (SVM-DUAL) then $w^* = w^*(x^*) = \frac{1}{\lambda m} \sum_{i=1}^m y^{(i)} (x^*)^{(i)} A_i^T$ is an optimal solution of the primal problem

$$\min_{w \in \mathbf{R}^N} P(w) := \frac{1}{N} \sum_{i=1}^N \mathcal{L}(w, A_i, y^{(i)}) + \frac{\lambda}{2} \|w\|^2, \quad (27)$$

where $\mathcal{L}(w, A_i, y^{(i)}) = \max\{0, 1 - y^{(i)} A_i w\}$.

Our second example implementation is a distributed coordinate-descent algorithm for support vector machines (SVM) in the (SVM-DUAL) formulation. In this case, we define

$$g_k := \frac{1}{\lambda m} \sum_{i=1}^m x_k^{(i)} y^{(i)} A_i^T. \quad (28)$$

Then

$$\nabla_i f(x) = \frac{y^{(i)} A_i g_k - 1}{m}, \quad L_i = \frac{\|A_i\|^2}{\lambda m^2}. \quad (29)$$

The optimal step length is then solution of a one-dimensional problem:

$$h^{(i)}(x_k) = \arg \min_{t \in \mathbf{R}} \nabla_i f(\alpha) t + \frac{\beta}{2} L_i t^2 + \Phi_{[0,1]}(\alpha^{(i)} + t) \quad (30)$$

$$= \text{clip}_{[-\alpha^{(i)}, 1-\alpha^{(i)}]} \left(\frac{\lambda m (1 - y^{(i)} A_i g_k)}{\beta \|A_i\|^2} \right), \quad (31)$$

where for $a < b$

$$\text{clip}_{[a,b]}(\zeta) = \begin{cases} a, & \text{if } \zeta < a, \\ b, & \text{if } \zeta > b, \\ \zeta, & \text{otherwise.} \end{cases}$$

The new value of the auxiliary vector $g_{k+1} = g(x_{k+1})$ is given by

$$g_{k+1} = g_k + \underbrace{\sum_{c=1}^C \sum_{i \in Z_k^{(c)}} \frac{1}{\lambda m} h^{(i)}(x_k) y^{(i)} A_i^T}_{\delta g^{(c)}} \quad (32)$$

and the duality gap $G(x_k) = P(g_k) + F(x_k)$ can be easily obtained [36, 35, 37] as

$$G(x_k) = \frac{1}{m} \sum_{i=1}^m (\mathcal{L}(g_k, A_{i:}, y^{(i)}) - x_k^{(i)}) + \lambda \|g_k\|^2. \quad (33)$$

7 Per-iteration complexity

Using to the auxiliary vector g_k , which was introduced in the previous section, Algorithm 1 has two alternating and time consuming sub-procedures, namely:

1. computation of an update $\sum_{i \in Z_k^{(c)}} U_i h^{(i)}(x_k)$ and the accumulation of g_k : $\delta g^{(c)}$,
2. updating g_k to g_{k+1} .

Let us denote the run-time of the first sub-procedure by $\mathcal{T}_1(\tau)$, considering this depends on τ , and the run-time of a second one by \mathcal{T}_2 . We will neglect the rest of the run-time cost, such as managing a loop, evaluation of termination criteria, measuring a computation time, etc. The total run-time cost \mathcal{T}_T is hence given by

$$\mathcal{T}_T = \mathcal{O} \left(\frac{\beta}{C\tau} (\mathcal{T}_1(\tau) + \mathcal{T}_2) \right) \quad (34)$$

where we consider the case when $\mu_\Omega(w) \equiv 0$ in (21). Let us now for simplicity assume that the first sub-procedure is linear in τ , i.e., $\mathcal{T}_1(\tau) = \tau \mathcal{T}_1(1) =: \tau \mathcal{T}_1$. Then

$$\mathcal{T}_T = \mathcal{O} \left(\frac{\beta}{C\tau} (\tau \mathcal{T}_1 + \mathcal{T}_2) \right). \quad (35)$$

Numerical values of \mathcal{T}_1 and \mathcal{T}_2 could be estimated, given problem sparsity and underlying hardware, or can be measured during the run.

Optimal choice of sampling parameter τ . In the previous paragraph, we gave an estimate of the complexity of a single iteration. In this paragraph, we answer the question of how to choose a τ given times $\mathcal{T}_1, \mathcal{T}_2$. For variable β , we have more options, but we stick to the most general one given in (11). Given that $s \geq 2$, we have

$$\mathcal{T}_T = \mathcal{O} \left(\frac{1 + \frac{(\xi-1)(\tau-1)}{s-1} + (C-1) \frac{\xi\tau}{s}}{C} \left(r_{1,2} + \frac{1}{\tau} \right) \mathcal{T}_2 \right) = \mathcal{O} \left(\left(\frac{s}{\xi C} + \tau \right) \left(r_{1,2} + \frac{1}{\tau} \right) \right), \quad (36)$$

where $r_{1,2} = \frac{\mathcal{T}_1}{\mathcal{T}_2}$ is a work to communication ratio. The optimal parameter τ^* can be obtain by minimizing (36) and is given by

$$\tau^* = \sqrt{\frac{s}{r_{1,2} \xi C}}. \quad (37)$$




Fig. 2 An illustration of a naïve (PS) approach (left), which alternates between parallel regions, where computations take place, and serial regions dedicated to MPI communications with other computers. An alternative (FP) approach (right) dedicates the communication task to one thread and uses other threads for computation.

Therefore, smaller values of $r_{1,2}$ imply that we should do more work in each iteration, and hence bigger values of τ should be chosen. This is quite natural, as one should tune the parameters in such a way that time spent in communication should be in comparable with that of effective computation.

Message Passing Interface (MPI). In order to discuss finer details of the implementations, we need to introduce the architecture we use. We use OpenMP [38] for dealing with concurrency within a single computer and Message Passing Interface (MPI) [39] as the abstraction layer for network communication. In MPI, one passes data from one MPI process to another MPI process, which may run on another computer. (We disregard the concept of groups for brevity.) Communication can involve any subset of computers, which run MPI processes. Communication can be either blocking (“synchronous”) or non-blocking (“asynchronous”). A *collective* operation involves the communication among two or more MPI processes. An example of a collective operation is a *barrier*, where computers wait until all of them reach the same point in the algorithm. Another common collective operation is *reduce all*, which is parametrized by an arbitrary operation that takes a set of elements and produces a single element of the same type. This “reduce” operation is applied to all elements of the particular type stored across all MPI processes and the result is returned to all MPI processes. For example, let us assume that each computer stores a vector $\delta g^{(c)} \in \mathbf{R}^m$ and the goal is to sum it up, i.e., to compute $\delta g^{(1,\dots,C)} = \sum_{c=1}^C \delta g^{(c)}$ and to make this result available on each computer. Figure 3 shows a standard approach, which leads to the desired result. From the performance point of view, however, the use of *reduce all* should be minimized, as it involves an implicit synchronisation and leaves most of the computers idle throughout the collective operation.

This suggests the following range of progressively better-performing variants:

Alternating Parallel and Serial regions (PS). The naïve implementation alternates two sub-procedures. One, which is computationally heavy and is done in parallel, but with no MPI communication, and another one, which is purely communicational. As an easy fix, one can dedicate one thread to the communication and other threads within the same computer to computation. We call this approach **Fully Parallel (FP)**. Figure 2 compares the naïve strategy (left) with the FP (right),

Reduce All (RA). As mentioned above, the use of *reduce all* operations significantly decreases the performance of many distributed algorithms. It is, however, the preferred form of communication between computers close to each other in the





Fig. 3 Schematic diagram of a standard reduce all implementation. The goal is to compute $\sum_{c=1}^C \delta g^{(c)}$. The arrows show data flow between computers.

Fig. 4 Illustration of ASL method for $C = 4$. During k -th iteration, computer c obtains its contribution $\delta g_k^{(c)}$ but asynchronously sends an accumulated update $\delta G_k^{(c)}$ to its successor.



computer network, such as computers directly connected by a network cable. The use of asynchronous methods is also preferred over synchronous methods.

Asynchronous StreamLined (ASL). We propose another pattern of communication, where each computer in one iteration sends only one message to the closest computer, asynchronously, and receives only one message from another computer close-by, asynchronously. The communication hence takes place in a ring. This tweak, however, requires a significant change in the algorithm. Figure 4 illustrates the data flow of messages at the end of iteration k for $C = 4$. We fix an order of computers in a ring, denoting $\text{pred}_R(c)$ and $\text{succ}_R(c)$ the two computers neighbouring computer c along the two directions on the ring. Computer c always receives data only from computer $\text{pred}_R(c)$ and sends data only to computer $\text{succ}_R(c)$. Let us denote by $\delta G_k^{(c)}$ the data, which computer c sends to computer $\text{succ}_R(c)$ at the end of iteration k . When computer c starts iteration k , it has already received $\delta G_{k-1}^{(\text{pred}_R(c))}$ ³. Hence the data, which will be sent at the end of iteration k by computer c are:

$$\delta G_k^{(c)} = \delta G_{k-1}^{(\text{pred}_R(c))} - \delta g_{k-C}^{(c)} + \delta g_k^{(c)}. \quad (38)$$

It should be noticed that at the end of each iteration in the ASL procedure, each computer has a different vector g_k , which we denote $g_k^{(c)}$. The update rule is

$$g_{k+1}^{(c)} = g_k^{(c)} + \delta g_k^{(c)} + \delta G_k^{(\text{pred}_R(c))} - \delta g_{k-C+1}^{(c)}. \quad (39)$$

The clear advantage of the ASL method is a decrease in communication time. On the other hand it comes with a cost of slower propagation of information. Indeed, it

³ For the start of the algorithm we define $\delta g_l^{(c)} = \delta G_l^{(c)} = \mathbf{0}$ for all $l < 0$.

Table 2 Summary of additional memory and computation requirements for strategies RA, SLA, AST.

| strategy | memory for g 's | communication | extra computation |
|----------|-------------------|--|-------------------|
| RA | $2m$ | \mathcal{T}_{ra} | 0 |
| SLA | $(2 + C)m$ | \mathcal{T}_{p2p} | $4m$ additions |
| AST | $(2 + C/r)m$ | $\mathcal{T}_{p2p} + \mathcal{T}_{ra}/r$ | $8m$ additions |

takes $C - 1$ iterations to propagate information to all computers. It also comes with bigger storage requirements, as at iteration k , we have to have all vectors $\delta g_l^{(c)}$ for $k - C \leq l \leq k$ stored on computer c .

Asynchronous Torus (AST). There is a compromise solution, though, which inherits many desirable features of both RA and ASL. This employs a toroidal networking topology, which is common in high-performance computing (HPC) in general, and HPC using InfiniBand networks [40], in particular. Let us assume that C is a multiple of $r \in \mathbb{N}$, where r represents the width of a torus, i.e., C computers are partitioned into subsets R_i each with size r . Each group R_i has a root computer. These root computers aggregate updates from their respective groups, e.g., using a local *reduce all* operation, in each iteration and exchange those update in an asynchronous ring with two other adjacent root computers. Thus the communication between the root nodes follows the ASL communication pattern. The AST approach decreases the propagation time from C to $\frac{C}{r}$, additional storage is also decrease by factor r , and the overall communication complexity remains low.

The Comparison. Changing from the FP approach to the PS approach does not require much computational or storage overhead, but can reduce the idle time of processors. However, changing from RA to SLA or AST brings significant storage requirements, while it reduces both communication and idle time significantly. Table 2 summarize maximum memory requirements on each single node of the cluster, time spent in communication, and amount of data transferred over the network. Once the time spent in communication is measured or estimated, one can pick the most appropriate strategy. Notice that the wall-clock time required for the *reduce all* operation, \mathcal{T}_{ra} , is typically of the order $\mathcal{O}(\log C) \cdot \mathcal{T}_{p2p}$, where \mathcal{T}_{p2p} is the time required by the point-to-point transmission.

8 Numerical experiments

In this section we present numerical evidence of the efficiency of the distributed (block) coordinate-descent method.

The code. The code of the distributed (block) coordinate-descent solver is part of our AC-DC library, available at <http://code.google.com/p/ac-dc/>. The library is written in C++ using OpenMP. The extensive use of template classes, Boost::MPI, and Boost.Serialization makes it easy to change the composite function and the precision of the computation. Both wall-clock and CPU-time were measured




Fig. 5 Evolution of duality gap for the WebSpam dataset for various choices of τ .

using Boost::Timers, which achieve nano-second accuracy on recent processors running recent versions of Linux.

The facility. Our empirical tests were conducted in UK's high-performance computing facility, HECToR, equipped with multi-core computers connected using Infiniband [40]. In particular, in Phase 3 of the facility, which is a Cray XE6 cluster, we have used up to 128 nodes, equipped with two AMD Opteron Interlagos 16-core processors and 32 GB of memory each. This gave us 4,096 cores in total, interconnected using Cray Gemini routers in a 3D torus. Each Gemini router was connected to processors and random-access memory of two nodes via HyperTransport links. Each router is then connected to ten other routers. In practice, the latency is about 1–1.5 microseconds and the capacity of each link is 8 GBs⁻¹. The facility ran a Cray Linux Environment, based on SuSE Linux.

Support Vector Machines (SVM). One of the goals of this paper is to train huge sparse support-vector machines (SVM) that do not fit into the memory of a single computer. In the machine learning literature, one often performs experiments on instances of moderate size, e.g., 100 MB [41, 25, 36]. Well-known instances of this scale include, e.g., CCAT variant of RCV1 [42], Astro-ph [41], and COV [41]. In this Section, we focus on a larger dataset, known as WebSpam [43]. This dataset consists of 350,000 observations (rows) and 16,609,143 features (columns). The size of the instance is 25 GB. Figure 5 show the execution time and duality gap for WebSpam dataset, using $C = 16$ MPI processes, with each process using 8 threads. τ is the number of coordinates updated by one MPI process during one iteration. As expected, the main run-time cost is not computing the updates, but updating g . Let us remark that ϵ is usually not particularly small in the machine-learning community. In experimenting with small ϵ , we just wanted to demonstrate that our algorithm is able to close the duality gap within the limits of machine precision. The truly important measures of the performance of the classifier, e.g., 0-1 loss or prediction error, are actually within 10 % after the first minute, which is the first time we compute it. In practice, a duality gap of 0.1 or 0.01 can be sufficient for machine learning problems.

Sparse least squares (LASSO). Next, we solved an artificial instance of sparse least squares with a matrix of $n = 10^9$ rows and $d = 5 \cdot 10^8$ columns in block-angular form:

$$A = \begin{pmatrix} A_{loc}^{(1)} & 0 & \cdots & 0 \\ 0 & A_{loc}^{(2)} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ A_{glob}^{(1)} & A_{glob}^{(2)} & \cdots & A_{glob}^{(C)} \end{pmatrix}. \quad (40)$$

requiring 3 TB to store. Such matrices often arise in stochastic optimization. We used 128 nodes with 4 MPI processes on each node. Each MPI process ran 8 OpenMP threads, giving a total of 4,096 hardware threads. Each node c stored two matrices: $A_{loc}^{(c)} \in \mathbf{R}^{1,952,148 \times 976,562}$ and $A_{glob}^{(c)} \in \mathbf{R}^{500,224 \times 976,562}$. The average number of non-zero elements per row is 175 and 1,000 for $A_{loc}^{(c)}$ and $A_{glob}^{(c)}$, respectively. When communicating $g_k^{(c)}$, only entries corresponding to the global part of $A^{(c)}$ need to be communicated, and hence in RA, a *reduce all* operation is applied to vectors $\delta g_{glob}^{(c)} \in \mathbf{R}^{500,224}$. In ASL, vectors with the same length are sent. The optimal solution x^* has exactly 160,000 nonzero elements. Figure 6 compares the evolution of $F(x_k) - F^*$ for ASL-FP and RA-FP.




Fig. 6 Evolution of $F(x_k) - F^*$ in time. ASL-FP significantly outperforms RA-FP. The loss F is pushed down by 25 degrees of magnitude in less than 30 minutes (3TB problem).

9 Conclusions

Overall, distributed algorithms can be both very efficient and easy to implement, when one picks the right approach. The first steps taken by the present authors over the past two years seem to have been validated by the considerable interest [26, 27, 37] they have generated.

References

1. P. Richtárik and M. Takáč, “Parallel coordinate descent methods for big data optimization,” *arXiv:1212.0873*, 2012.
2. D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1989.
3. O. Fercoq and P. Richtárik, “Accelerated, parallel and proximal coordinate descent,” *arXiv:1312.5799*, 2013.
4. A. Saha and A. Tewari, “On the finite time convergence of cyclic coordinate descent methods,” *SIAM J. Optimiz.*, vol. 23, no. 1, pp. 576–601, 2013.
5. P. Richtárik and M. Takáč, “Efficient serial and parallel coordinate descent methods for huge-scale truss topology design,” in *Operations Research Proceedings 2011*, pp. 27–32, Springer, 2012.
6. Y. Nesterov, “Efficiency of coordinate descent methods on huge-scale optimization problems,” *SIAM J. Optimiz.*, vol. 22, no. 2, pp. 341–362, 2012.
7. P. Richtárik and M. Takáč, “Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function,” *Math. Program.*, vol. 144, no. 1-2, pp. 1–38, 2014.
8. O. Fercoq and P. Richtárik, “Smooth minimization of nonsmooth functions with parallel coordinate descent methods,” *arXiv:1309.5885*, 2013.
9. Z. Lu and L. Xiao, “On the complexity analysis of randomized block-coordinate descent methods,” *arXiv preprint arXiv:1305.4723*, 2013.
10. I. Necoara and D. Clipici, “Distributed coordinate descent methods for composite minimization,” *arXiv:1312.5302*, 2013.
11. Y. T. Lee and A. Sidford, “Efficient accelerated coordinate descent methods and faster algorithms for solving linear systems,” in *54th Annual Symposium on Foundations of Computer Science*, IEEE, 2013.
12. D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, 2nd ed., Sept. 1999.
13. Z. Q. Luo and P. Tseng, “A coordinate gradient descent method for nonsmooth separable minimization,” *J. Optim. Theory Appl.*, vol. 72, January 2002.
14. P. Tseng and S. Yun, “A coordinate gradient descent method for nonsmooth separable minimization,” *Math. Program.*, vol. 117, pp. 387–423, July 2008.
15. P. Tseng and S. Yun, “Block-coordinate gradient descent method for linearly constrained nonsmooth separable optimization,” *J. Optim. Theory Appl.*, vol. 140, pp. 513–535, 2009.
16. P. Tseng, “Convergence of a block coordinate descent method for nondifferentiable minimization,” *J. Optim. Theory Appl.*, vol. 109, no. 3, pp. 475–494, 2001.
17. R. Tappenden, P. Richtárik, and M. Takáč, “Improved complexity analysis of parallel coordinate descent methods,” 2014. Technical Report, the University of Edinburgh.
18. R. Tappenden, P. Richtárik, and J. Gondzio, “Inexact coordinate descent: complexity and preconditioning,” *arXiv:1304.5530*, 2013.
19. J. Liu, S. J. Wright, C. Ré, and V. Bittorf, “An asynchronous parallel stochastic coordinate descent algorithm,” *arXiv:1311.1873*, 2013.

20. P. Richtárik and M. Takáč, “On optimal probabilities in stochastic coordinate descent methods,” *arXiv:1310.3438*, 2013.
21. C. Scherrer, A. Tewari, M. Halappanavar, and D. Haglin, “Feature clustering for accelerating parallel coordinate descent,” *Advances in Neural Information Processing Systems*, vol. 25, pp. 28–36, 2012.
22. R. Tappenden, P. Richtárik, and B. Büke, “Separable approximations and decomposition methods for the augmented lagrangian,” *Optim. Method. Softw.*, 2014. *arXiv:1308.6774*.
23. P. Zhao and T. Zhang, “Stochastic optimization with importance sampling,” *arXiv:1401.2753*, 2014.
24. F. Niu, B. Recht, C. Ré, and S. J. Wright, “Hogwild!: A lock-free approach to parallelizing stochastic gradient descent,” *Advances in Neural Information Processing Systems*, vol. 24, pp. 693–701, 2011.
25. C.-J. Hsieh, K.-W. Chang, C.-J. Lin, S. S. Keerthi, and S. Sundararajan, “A dual coordinate descent method for large-scale linear SVM,” in *Proceedings of the 25th International Conference on Machine Learning, ICML ’08*, (New York, NY, USA), pp. 408–415, ACM, 2008.
26. P. Richtárik and M. Takáč, “Distributed coordinate descent method for learning with big data,” *arXiv:1310.2059*, 2013.
27. O. Fercoq, Z. Qu, P. Richtárik, and M. Takáč, “Fast distributed coordinate descent for non-strongly convex losses,” *IEEE Workshop on Machine Learning for Signal Processing*, 2014.
28. N. S. M. Salleh, A. Suliman, and A. R. Ahmad, “Parallel execution of distributed SVM using MPI (CoDLib),” in *Information Technology and Multimedia (ICIM)*, pp. 1–4, IEEE, 2011.
29. E. Y. Chang, K. Zhu, H. Wang, H. Bai, J. Li, Z. Qiu, and H. Cui, “PSVM: Parallelizing support vector machines on distributed computers,” *Advances in Neural Information Processing Systems*, vol. 20, 2007.
30. N. K. Alham, M. Li, Y. Liu, and S. Hammoud, “A MapReduce-based distributed SVM algorithm for automatic image annotation,” *Comput. Math. Appl.*, vol. 62, no. 7, pp. 2801 – 2811, 2011.
31. Y. Nesterov, *Introductory lectures on convex optimization*, vol. 87 of *Applied Optimization*. Kluwer, 2004.
32. B. K. Natarajan, “Sparse approximate solutions to linear systems,” *SIAM journal on computing*, vol. 24, no. 2, pp. 227–234, 1995.
33. D. Ge, X. Jiang, and Y. Ye, “A note on the complexity of ℓ_p minimization,” *Math. Program.*, vol. 129, no. 2, pp. 285–299, 2011.
34. A. Patrascu and I. Necoara, “Random coordinate descent methods for ℓ_0 regularized convex optimization,” *arXiv:1403.6622*, 2014.
35. S. Shalev-Shwartz and T. Zhang, “Stochastic dual coordinate ascent methods for regularized loss,” *J. Mach. Learn. Res.*, vol. 14, no. 1, pp. 567–599, 2013.
36. M. Takáč, A. S. Bijral, P. Richtárik, and N. Srebro, “Mini-batch primal and dual methods for SVMs,” *J. Mach. Learn. Res.*, vol. W&CP 28, pp. 1022–1030, 2013.
37. M. Jaggi, V. Smith, M. Takáč, J. Terhorst, T. Hofmann, and M. I. Jordan, “Communication-efficient distributed dual coordinate ascent,” *Advances in Neural Information Processing Systems*, vol. 27, 2014.
38. OpenMP Architecture Review Board, *OpenMP Application Program Interface*. 2011.
39. M. Snir, S. Otto, S. Huss-Lederman, D. Walker, and J. Dongarra, *MPI-The Complete Reference, Volume 1: The MPI Core*. Cambridge, MA, USA: MIT Press, 2nd. (revised) ed., 1998.
40. InfiniBand Trade Association, *InfiniBand Architecture Specification, Volume 1, Release 1.0*. 2005.
41. S. Shalev-Shwartz, Y. Singer, N. Srebro, and A. Cotter, “Pegasos: Primal estimated sub-gradient solver for SVM,” *Math. Program.*, vol. 127, no. 1, pp. 3–30, 2011.
42. D. D. Lewis, Y. Yang, T. G. Rose, and F. Li, “Rcv1: A new benchmark collection for text categorization research,” *J. Mach. Learn. Res.*, vol. 5, pp. 361–397, Dec. 2004.
43. L. Data, 25/9/2014. <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>.

Notation Glossary

Optimization problem

| | | |
|----------|---|-----|
| N | dimension of the optimization variable | (1) |
| x, h | vectors in \mathbf{R}^N | |
| F | $F = f + \Omega$ (loss / objective function) | (1) |
| F^* | optimal value, we assume $F^* > -\infty$ | |
| f | smooth convex function ($f : \mathbf{R}^N \rightarrow \mathbf{R}$) | (1) |
| Ω | convex block separable function ($\Omega : \mathbf{R}^N \rightarrow \mathbf{R} \cup \{+\infty\}$) | (1) |

Block structure

| | | |
|-----------------|---|-------|
| n | number of blocks | |
| $[n]$ | $[n] = \{1, 2, \dots, n\}$ (the set of blocks) | Sec 2 |
| N_i | dimension of block i ($N_1 + \dots + N_n = N$) | Sec 2 |
| U_i | an $N_i \times N$ column submatrix of the $N \times N$ identity matrix | Sec 2 |
| $x^{(i)}$ | $x^{(i)} = U_i^T x \in \mathbf{R}^{N_i}$ (block i of vector x) | Sec 2 |
| $\nabla_i f(x)$ | $\nabla_i f(x) = U_i^T \nabla f(x)$ (block gradient of f associated with block i) | Sec 2 |
| L_i | block Lipschitz constant of the gradient of f | (5) |
| L | $L = (L_1, \dots, L_n)^T \in \mathbf{R}^n$ (vector of block Lipschitz constants) | |
| w | $w = (w_1, \dots, w_n)^T \in \mathbf{R}^n$ (vector of positive weights) | |
| $\ x\ _w$ | $\ x\ _w = (\sum_{i=1}^n w_i \ x^{(i)}\ _{(i)}^2)^{1/2}$ (weighted norm associated with x) | (3) |
| Ω_i | i -th componet of $\Omega = \Psi_1 + \dots + \Omega_n$ | (7) |
| $\mu_\Omega(W)$ | strong convexity constant of Ω with respect to the norm $\ \cdot\ _w$ | (20) |
| $\mu_f(W)$ | strong convexity constant of f with respect to the norm $\ \cdot\ _w$ | (20) |
| J | subset of $\{1, 2, \dots, n\}$ | |
| $x_{[Z]}$ | vector in \mathbf{R}^N formed from x by zeroing out blocks $x^{(i)}$ for $i \notin Z$ | (2) |

Block samplings

| | | |
|-----------------------|---|--------------|
| ω | degree of partial separability of f | Assumption 1 |
| \hat{Z}, Z_k | distributed block samplings (random subsets of $\{1, 2, \dots, n\}$) | Sec 3 |
| C | number of nodes (partitions) | Sec 3 |
| τ | # of blocks updated in 1 iteration within one partition | |
| $\{P^{(c)}\}_{c=1}^C$ | partition of $[n]$ onto C parts | |

Algorithm

| | | |
|--------------|---|-----|
| β | stepsize parameter depending on f and \hat{Z} | |
| $h^{(i)}(x)$ | $h^{(i)}(x) = (h(x))^{(i)} = \arg \min_{t \in \mathbf{R}^{N_i}} \langle \nabla_i f(x), t \rangle + \frac{\beta w_i}{2} \ t\ _{(i)}^2 + \Omega_i(x^{(i)} + t)$ | (8) |